

Sensor Selection for Event Detection in Wireless Sensor Networks

Dragana Bajović, Bruno Sinopoli and João Xavier

Abstract

We consider the problem of sensor selection for event detection in wireless sensor networks (WSNs). We want to choose a subset of p out of n sensors that yields the best detection performance. As the sensor selection optimality criteria, we propose the Kullback-Leibler and Chernoff distances between the distributions of the selected measurements under the two hypothesis. We formulate the maxmin robust sensor selection problem to cope with the uncertainties in distribution means. We prove that the sensor selection problem is NP hard, for both Kullback-Leibler and Chernoff criteria. To (sub)optimally solve the sensor selection problem, we propose an algorithm of affordable complexity. Extensive numerical simulations on moderate size problem instances (when the optimum by exhaustive search is feasible to compute) demonstrate the algorithm's near optimality in a very large portion of problem instances. For larger problems, extensive simulations demonstrate that our algorithm outperforms random searches, once an upper bound on computational time is set. We corroborate numerically the validity of the Kullback-Leibler and Chernoff sensor selection criteria, by showing that they lead to sensor selections nearly optimal both in the Neyman-Pearson and Bayes sense.

Keywords: sensor selection, event detection, robust optimization, Chernoff distance, Kullback-Leibler distance

Partially supported by grants SIPM PTDC/EEA-ACR/73749/2006 and SFRH/BD/33517/2008 (through the Carnegie Mellon/Portugal Program managed by ICTI) from Fundação para a Ciência e Tecnologia and also by ISR/IST plurianual funding (POSC program, FEDER).

D. Bajović is with the Institute for Systems and Robotics (ISR), Instituto Superior Técnico (IST), Lisbon, Portugal, and with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA dragana@isr.ist.utl.pt, dbajovic@andrew.cmu.edu

B. Sinopoli is with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA brunos@ece.cmu.edu

J. Xavier is with the Institute for Systems and Robotics (ISR), Instituto Superior Técnico (IST), Lisbon, Portugal jxavier@isr.ist.utl.pt

I. INTRODUCTION

Wireless sensor networks (WSNs) operate with limited power and communication resources. When observing phenomena with WSNs, a major challenge is to balance the tradeoff between the quality and the cost of operation. A fundamental problem of this kind, the sensor selection problem, is how to optimally select a limited subset of sensors (hence limiting the operation cost) that gives the most valuable information about the observed phenomena.

Problem statement. This paper studies the sensor selection problem for event detection in WSNs. Nature can be in one of two states: H_1 (event occurring, e.g., target present) and H_0 (event not occurring, e.g., target absent). A WSN, composed of n sensors, instruments the nature. The distribution of the n -dimensional measurement vector is assumed Gaussian under the two hypothesis, with different means m_i and covariances S_i , $i = 0, 1$. We assume that, due to inherent WSN constraints, such as power, only p (out of n) sensors can sense and transmit their readings to a fusion node; based on the received p readings, the fusion node performs detection, i.e., decides which of the two hypothesis is true. We ask the following question: *Which p sensors should be chosen to achieve the best detection performance?*

Each possible p -sensor selection induces a p dimensional Gaussian distribution π_i (of selected sensors) under hypothesis H_i , $i = 0, 1$. Intuitively, p -sensor selection that yields more distant distributions π_1 and π_0 leads to better detection. Hence, we propose, as sensor selection optimality criteria: 1) the Kullback-Leibler (KL) distance; and 2) the Chernoff (C) distance between π_1 and π_0 . In practice, the distribution parameters (m_i, S_i) are estimated from training data, and may not be known exactly, but only within an uncertainty region. We thus formulate the robust maxmin sensor selection problem of maximizing the KL (or C) distance between the selected distributions π_1 and π_0 , for the worst case of parameter drifts. In this paper, we address the case when only the means of the two distributions are uncertain, relegating the general case for future work.

Contributions. The problem of evaluating the best p -sensor subset is combinatorial; checking over all $\binom{n}{p}$ possible combinations becomes infeasible when n and p are sufficiently large. We indeed prove that the KL and C sensor selection problems are NP hard; hence, it is unlikely to find an efficient algorithm that solves large instances of these problems. To (suboptimally in general) solve the sensor selection problems, we propose a computationally affordable algorithm. For example, to select 10 out of 100 sensors, our algorithm takes only few seconds on a current generation personal computer. There is no theoretical guarantee that our algorithm produces an optimal or near optimal solution; however, extensive numerical experiments demonstrate that our algorithm produces an optimal (or near optimal) solution in most cases.

The KL and C distances are only heuristic optimality measures for the detection performance; the exact (yet in our case intractable) criterion is the probability of detection error (either in Neyman-Pearson or Bayes sense.) However, interestingly enough, we show by simulations that optimizing the KL and C distances yields to sensor selections that are very close to the optimal probability of error (either in Bayes or Neyman-Pearson sense). Moreover, they are often indeed optimal.

Selecting p out of n sensors is equivalent to finding the $n \times p$ selection matrix with 0/1 entries that maps the n -dimensional vector x of all measurements to the p -dimensional vector y of selected measurements. Our methodology for solving the combinatorial sensor selection problems relies on enlarging the search space from the 0/1 selection matrices to the set of Stiefel matrices (the matrices with orthonormal columns). Then, after solving for the Stiefel matrix, we project it back to the set of 0/1 selection matrices. The relaxed Stiefel problem corresponds to finding the linear map $\mathbb{R}^n \rightarrow \mathbb{R}^p$ which maximizes the (worst-case) KL (C) distance between the projected distributions in the lower p -dimensional space. To our best knowledge, existing work on this topic either does not solve the problem in full generality (i.e., unequal means and covariance matrices) or does not guarantee global optimality of their solutions (see [1], [2] for problems involving Chernoff distance and the closely related J-divergence). A major contribution of this paper is that we solve this nonconvex problem globally for the case $p = 1$, and in full generality, by reducing it to scalar (1D) problem over a compact interval. We tackle the generic case $p > 1$ via an incremental, greedy approach, based on the 1D case, which provides near optimal result with small computational cost.

This paper is related to our prior work [3], [4], which also considers sensor selection for event detection, but only based on the KL distance. With respect to KL distance, this paper provides a new heuristic with reduced complexity; more importantly, this paper studies the problem with respect to the Chernoff distance, which we did not consider in [3], [4]. With respect to [3], [4], this paper also contributes by validating the KL and C distances as good optimality criteria by showing their near optimality in the error probability sense, and by establishing NP hardness of the corresponding sensor selection optimization problems.

Finally, we would like to note that the KL-based and C-based sensor selection problems could be, in principle, globally solved (inefficiently) by, e.g., branch and bound methods [5], [6]. However, the computational time of such methods is often very long, even for modest values of n and p . We discuss in more detail the challenges to solve the sensor selection problems that we address at the end of subsection II-B.

Review of existing work on sensor selection. Sensor selection problems have been extensively studied in

different contexts, including target tracking (e.g., [7]), target localization (e.g., [8]), robotics (e.g., [9]), and wireless sensor networks (e.g. [10], [11]). Generally, all their work aims to optimize certain measure of performance of the system (e.g., utility [12], information theoretic measure [7], [13], estimation error [14]) subject to energy constraints (e.g., limited number of sensors to be selected [14]). Sensor selection can be geometric-based (e.g., [15]), or information-theoretic based (e.g., [7], [13]). Our work belongs to the class of sensor selection problems for inference tasks, i.e., sensor selection for estimation and detection, and it adopts the information-theoretic point of view. Reference [14] considers the problem of sensor selection for parameter estimation in WSNs, proposing to select the subset of p (out of n) sensors that minimize the determinant of the estimator covariance matrix. Reference [16] proposes distributed algorithms to (suboptimally) solve the sensor selection problem formulated in [14]. Reference [17] addresses the problem of selecting the maximal number of reliable sensors for estimation. Reference [18] shows, through the optimal experiment design framework ([19]) and using convex analysis, that optimal estimation is achievable by using only a relatively small number of sensors.

Paper organization. Section II introduces the model and formulates the sensor selection optimization problems. Section III details the algorithms for solving the robust sensor selection problems, in the presence of uncertainties in the means of the distributions. Section IV considers the special, yet important case, when there are no uncertainties in the distribution parameters. Section V demonstrates numerically that the KL and the C distances are appropriate metrics for sensor selection. Section VI shows near optimal performance of the proposed algorithms. Finally, Section VII concludes the paper.

II. PROBLEM MODEL: FORMULATION OF THE SENSOR SELECTION OPTIMIZATION PROBLEM

A. Problem model

We assume that nature can be in one of two states: H_1 —event occurring, and H_0 —event not occurring. Let $x \in \mathbb{R}^n$ denote the vector that collects all sensor measurements (one measurement per sensor). We assume that x is Gaussian under the hypothesis H_1 (respectively, H_0), with generic means and covariances (m_1, S_1) (respectively, (m_0, S_0)), i.e.,

$$\begin{aligned} H_0 &: x \sim \mathcal{N}(m_0, S_0) \\ H_1 &: x \sim \mathcal{N}(m_1, S_1) \end{aligned},$$

where $\mathcal{N}(\mu, \Sigma)$ denotes Gaussian distribution with the mean vector μ and the covariance matrix Σ . The Gaussian assumption on x is standard and can be, in many applications, justified, e.g, by central limit theorem arguments, see, e.g., [20], [21]. Noise correlation (i.e., non diagonal covariance matrices) is important to take into account in dense deployments of WSNs. We note that our formulation allows for

different covariances under two hypothesis ($S_1 \neq S_0$); in many applications, e.g., power-based detection of primary users for cognitive radios (see [21] for details), accounting for different covariances is essential.¹

Sensor selection. Sensors transmit their measurements to a fusion node, which conducts the hypothesis test. Due to power constraints, only p sensors, $p < n$, perform their measurements and transmit them to the fusion node. We address the problem of selecting the p sensors that guarantee the best detection performance. Mathematically, selecting p out of n sensors can be represented by a linear map $\mathbb{R}^n \rightarrow \mathbb{R}^p$, $x \mapsto y = E^\top x$, where $E \in \mathbb{R}^{n \times p}$ is a rank- p matrix that has exactly one unit entry per column, corresponding to a chosen sensor, and the other entries in columns being zero. The columns of E are orthonormal, i.e., $E^\top E = I_p$, where I_p denotes the $p \times p$ identity matrix. We refer to matrix E as the sensor selection matrix.

Hypothesis test induced by E . Conditioned on H_i , $i = 0, 1$, y is a linear transformation of a Gaussian vector x . Thus, y , under H_i , has the following distributions:

$$\begin{aligned} H_0 &: y \sim \mathcal{N}(E^\top m_0, E^\top S_0 E) \\ H_1 &: y \sim \mathcal{N}(E^\top m_1, E^\top S_1 E). \end{aligned} \quad (1)$$

The fusion node performs the following log-likelihood ratio (LLR) decision test:

$$\log \frac{f_1(y; E)}{f_0(y; E)} \underset{H_0}{\overset{H_1}{\gtrless}} \gamma, \quad (2)$$

where $f_i(\cdot; E)$, $i = 0, 1$, is the density of $\mathcal{N}(E^\top m_i, E^\top S_i E)$ and $\gamma \in \mathbb{R}$ is the test threshold, [2].

B. Formulation of the sensor selection optimization problem

Sensor selection optimality criteria. Detection performance is, as noted above, generally quantified by the error probabilities, P_{FA} , P_D and (P_e) . However, in the problem that we consider, none of the probabilities above admit closed form expression, and their minimization with respect to sensor selection is a very hard problem. As sensor selection optimality criteria, we propose the Kullback-Leibler (KL) and the Chernoff (C) distance between the tested distributions. Given two distributions, with densities f_1 and f_0 , KL and C distances measure dissimilarity between f_1 and f_0 , and they are defined as follows:

$$\begin{aligned} D_{KL}(f_1 \| f_0) &:= \int \log \frac{f_1(x)}{f_0(x)} f_1(x) dx \\ D_C(f_1, f_0) &:= \max_{s \in [0,1]} -\log \int f_1^s(x) f_0^{1-s}(x) dx. \end{aligned}$$

¹Generally, our model applies also if x contains samples from multiple sensors over *multiple sample times*, that is, several entries in x correspond to same sensor. In such scenario vector x accounts both for temporal and spatial correlation between observations. To focus the presentation, however, we assume that different entries of x correspond to different sensors.

Thus, our goal is to find a sensor selection that yields the maximal KL or C distance between the projected Gaussian distributions, $\mathcal{N}(E^\top m_1, E^\top S_1 E)$ and $\mathcal{N}(E^\top m_0, E^\top S_0 E)$ (see (1)).

We are motivated to choose KL and C distances, as sensor selection optimality criteria, by two fundamental results from detection theory: Chernoff-Stein lemma and Chernoff lemma. Chernoff-Stein lemma (resp. Chernoff lemma) states that, when the number of independent identically distributed (i.i.d.) observations grows, the rate of exponential decay of probability of false alarm, P_{FA} , (resp. probability of error, P_e) of the Neyman-Pearson optimal (resp. Bayes optimal) test equals the KL (resp. C) between the two distributions. Thus, for large number of samples, more distant distributions (in either KL or C sense) lead to better detection performance. Probabilistic distance measures have been often used in the literature as heuristics for detection problems (see, e.g., [1], [22] for applications in linear dimensionality reduction) and have shown excellent results, even when the number of samples is very small or even equal to one, see [22]. In section V, we demonstrate by numerical tests that the KL and C distances are indeed excellent criteria for sensor selection, exhibiting near optimal performance in the probability of error sense. Section V shows that, generally, C distance has an advantage in the regimes of high probability of detection (P_D) (upper part of the receiver operating characteristics (ROC) curve), while KL has an edge in the regimes of low P_{FA} (lower left part of the ROC curve).

Robustness against uncertainty in distribution means. We consider the case when the true distribution parameters (m_i, S_i) , $i = 0, 1$ are not exactly known at the fusion node detector (2). Fusion node has their estimates, (\hat{m}_i, \hat{S}_i) , which can be obtained, e.g., in the network training phase. Thus, there is a mismatch between the distribution used by the fusion node detector and the distribution that generates the observations. Our goal is to design a sensor selection that yields detection (2) robust to these mismatches. In this paper, we restrict our attention to the case where only the mean values are uncertain (the true covariance matrices are known), and we allow the mean values to drift in the following ellipsoidal uncertainty regions:

$$m_i \in \mathcal{E}(\hat{m}_i, k_i S_i^{-1}) \quad i = 0, 1. \quad (3)$$

Here \hat{m}_i denotes the estimated mean vector, $i = 0, 1$, $\mathcal{E}(a, A)$ (A is a positive definite matrix) denotes the ellipsoid

$$\mathcal{E}(a, A) = \{x \in \mathbb{R}^n : (x - a)^\top A(x - a) \leq 1\}.$$

and the parameter $k_i \in (0, +\infty]$ is a free parameter which controls the “size” of the uncertainty region, e.g, if $k_i = +\infty$, there is no uncertainty: $\hat{m}_i = m_i$. The orientations of the uncertainty ellipsoids in (3) are induced by the covariance matrices S_0 and S_1 . This choice of the form of uncertainty regions is

motivated by the following fact: if the means are estimated via the sample mean estimator based on N i.i.d. observations (the minimum variance unbiased estimate for Gaussian distributions), then the covariance of the estimate \hat{m}_i equals $\frac{1}{N}S_i$, for $i = 0, 1$. A standard measure of the uncertainty, the confidence region ([2]), is for m_i given by (3). The scaling constants k_0 and k_1 are, in this context, proportional to N ; also, if N is known, k_0 and k_1 can be used to design the uncertainty regions of desired confidence levels.

We address the uncertainties in the mean vectors adopting the worst case approach. That is, we search for the sensor selection that gives the maximal distance (KL and C) for the worst case of the mean parameter drift.

Optimization problems. We introduce the following two functions, f_{KL} and f_{C} , that capture the dependence of the KL and the C distance on the selection matrix E and the mean vectors m_0 and m_1 :

$$\begin{aligned} f_{\text{KL}}(E, m_0, m_1) &:= D_{\text{KL}}\left(\mathcal{N}(E^\top m_1, E^\top S_1 E) \parallel \mathcal{N}(E^\top m_0, E^\top S_0 E)\right) \\ \max_{s \in [0,1]} f_{\text{C}}(s, E, m_0, m_1) &:= D_{\text{C}}\left(\mathcal{N}(E^\top m_1, E^\top S_1 E), \mathcal{N}(E^\top m_0, E^\top S_0 E)\right). \end{aligned}$$

The robust sensor selection optimization problems are then given as follows:

$$\begin{aligned} \text{maximize} \quad & \min_{m_0 \in \mathcal{E}(\hat{m}_0, k_0 S_0^{-1}), m_1 \in \mathcal{E}(\hat{m}_1, k_1 S_1^{-1})} f_{\text{KL}}(E, m_0, m_1) \\ \text{subject to} \quad & E_{ij} \in \{0, 1\} \\ & E^\top E = I_p \end{aligned} \tag{4}$$

$$\begin{aligned} \text{maximize} \quad & \min_{m_0 \in \mathcal{E}(\hat{m}_0, k_0 S_0^{-1}), m_1 \in \mathcal{E}(\hat{m}_1, k_1 S_1^{-1})} \max_{s \in [0,1]} f_{\text{C}}(s, E, m_0, m_1) \\ \text{subject to} \quad & E_{ij} \in \{0, 1\} \\ & E^\top E = I_p \end{aligned} . \tag{5}$$

It can be shown that $\text{tr}(\cdot)$ and $|\cdot|$ denote the trace and the determinant, respectively):

$$f_{\text{KL}}(E, m_0, m_1) = \frac{1}{2} \left\{ (m_1 - m_0)^\top E \left(E^\top S_0 E \right)^{-1} E^\top (m_1 - m_0) + \text{tr} \left(\left(E^\top S_0 E \right)^{-1} E^\top S_1 E \right) - \log \frac{|E^\top S_1 E|}{|E^\top S_0 E|} - p \right\} \tag{6}$$

$$f_{\text{C}}(s, E, m_0, m_1) = \frac{1}{2} \left\{ s(1-s)(m_1 - m_0)^\top E \left(sE^\top S_0 E + (1-s)E^\top S_1 E \right)^{-1} E^\top (m_1 - m_0) - \log \frac{|E^\top S_0 E|^s |E^\top S_1 E|^{1-s}}{|sE^\top S_0 E + (1-s)E^\top S_1 E|} \right\}. \tag{7}$$

Optimization problems (4) and (5) are combinatorial. When n and p are small, a simple method for solving them is exhaustive search that checks all $\binom{n}{p}$ sensor subsets (i.e. selection matrices). For large n

and p , however, this method becomes computationally infeasible. Indeed, we have the following result, which we prove in the Appendix A:

Theorem 1 The optimization problems (4) and (5) are NP hard (even when $k_0 = k_1 = +\infty$ and $S_1 = S_0$).

In principle, besides exhaustive search, problems (4) and (5) can be (inefficiently) solved by branch and bound methods, see, e.g., [6]. However, complexity of these methods relies strongly on the choice of bounds on the cost function; finding tight bounds is a hard problem itself. An interesting method for solving a problem somewhat similar to ours is proposed in [23]. In [23], the authors address the problem of finding the most informative locations, for future sensor placements, in a discretized Gaussian field; the measure of informativeness that the authors propose is the mutual information between sensed and unsensed locations. The authors show that the mutual information is a submodular function, which assures that simple strategy of choosing sensor by sensor (greedy) gives a solution within $1 - \frac{1}{e}$ of the optimum. In our problems, however, such a bound on the greedy strategy (of selecting sensors one by one) does not hold, as our cost functions (KL and C distance) are not submodular (see the proof in Appendix B). In fact, in our problems, greedy performs poorly in many cases. The reason for this lies in the fact that the correlations—an information that greedy discards can play an important role in discerning between the two hypothesis.

III. SENSOR SELECTION ALGORITHMS

In this section, we present our algorithms, R-KL (robust KL based selection) and R-C (robust C-based selection), that, respectively, solve the problems (4) and (5). First, we explain the methodology and the structure of the algorithms, and we explain the geometrical intuition behind this methodology. Subsections III-A and III-B detail the R-KL and R-C algorithms.

Algorithms methodology and structure. Geometrically, one combination of p sensors defines one p -dimensional subspace of \mathbb{R}^n , spanned by a set of canonical basis vectors corresponding to the chosen sensors. We call this subspace a canonical subspace. The cost functions (6) and (7) depend on E only through its range, and problems (4) and (5), in this sense, search for the best subspace, among $\binom{n}{p}$ canonical subspaces, on which the original distributions should be projected. We relax these combinatorial problems by allowing for projections to arbitrary p -dimensional subspaces. Mathematically, this translates into replacing the set of 0/1 selection matrices with the set of $n \times p$ Stiefel matrices (that represent all p -dimensional subspaces). Then, we use a solution of the relaxed problem and “round” it by the closest canonical subspace.

We call the first phase of our algorithm, that solves the relaxed Stiefel problem, the *Relaxation phase*; the second phase, in which we find the closest canonical subspace, is the *Projection phase*. Finally,

the last, third step in our algorithms, the *Refinement phase*, refines the solution by performing local optimization.

A. Algorithm for the robust Kullback-Leibler based selection: the R-KL algorithm

1) *Relaxation phase:* We solve the following Stiefel relaxation of the problem (4):

$$\begin{aligned} & \text{maximize} \quad \min_{m_0 \in \mathcal{E}(\hat{m}_0, k_0 S_0^{-1}), m_1 \in \mathcal{E}(\hat{m}_1, k_1 S_1^{-1})} f_{\text{KL}}(E, m_0, m_1) \\ & \text{subject to} \quad E^\top E = I_p. \end{aligned} \quad (8)$$

The relaxed problem (8) is nonconvex and still difficult to solve. We globally solve this problem for the case $p = 1$, i.e., we find the best 1-dimensional projection. General, $p > 1$ case, is addressed by a greedy approach, using our $1 - D$ tool. We first detail the algorithm that solves $p = 1$ case.

(1.1) *Case $p = 1$: global solution* For $p = 1$, the constraint set of Stiefel matrices reduces to a sphere in \mathbb{R}^n and the problem (8) takes a simplified form:

$$\begin{aligned} & \text{maximize} \quad \min_{\substack{m_0 \in \mathcal{E}(\hat{m}_0, k_0 S_0^{-1}), \\ m_1 \in \mathcal{E}(\hat{m}_1, k_1 S_1^{-1})}} \frac{1}{2} \left\{ \frac{(e^\top (m_1 - m_0))^2}{e^\top S_0 e} + \frac{e^\top S_1 e}{e^\top S_0 e} - \log \frac{e^\top S_1 e}{e^\top S_0 e} - 1 \right\} \\ & \text{subject to} \quad e^\top e = 1 \end{aligned} \quad (9)$$

Our major contribution is showing that the problem (9) reduces to a search over a compact (one-dimensional) interval. We achieve this by a series of judicious problem reformulations, and by invoking convexity of quadratic mappings, [24], [25]. It is important to note that the original problem (9) has in general very high dimensionality (equal to the total number of sensor in the network n); also, due to its nonconvexity, it is very difficult to solve globally. By doing reformulations, we manage to map it to a tractable, scalar problem. The next Lemma, proved in [4], states the first step towards this goal. It shows that a solution of (9) can be reconstructed after solving a 2 dimensional problem (10).

Lemma 2 Suppose (x^*, y^*) solves

$$\begin{aligned} & \text{maximize} \quad \psi_{\text{KL}}(x, y) \\ & \text{subject to} \quad (x, y) \in \mathcal{R} \end{aligned} \quad (10)$$

where

$$\psi_{\text{KL}}(x, y) = x - \log x + \left\{ \left(\sqrt{y} - \frac{1}{\sqrt{k_1}} \sqrt{x} - \frac{1}{\sqrt{k_0}} \right)^+ \right\}^2, \quad (11)$$

$$\mathcal{R} = \left\{ (x, y) \in \mathbb{R}^2 : x = v^\top S v, y = v^\top M v, \text{ for some } v \in \mathbb{R}^n, v^\top v = 1 \right\}, \quad (12)$$

$x^+ = \max(0, x)$, $S = S_0^{-1/2} S_1 S_0^{-1/2}$, $m = S_0^{-1/2} (m_1 - m_0)$, $M = m m^\top$. Let $v^* \in \mathbb{R}^n$ be a unit-norm vector that generates x^* and y^* , i.e. $x^* = v^{*\top} S v^*$ and $y^* = v^{*\top} M v^*$. Then, $e^* := S_0^{-1/2} v^* / \|S_0^{-1/2} v^*\|$ solves (9).

Lemma 2 says that, in order to solve (9), it suffices to search over the set $\mathcal{R} \subset \mathbb{R}^2$. For $n \geq 3$, the set \mathcal{R} is compact and convex, as the image set of a unit sphere under two quadratic mappings, see [24]. Note that, since $\psi_{\text{KL}}(\cdot)$ is continuous and \mathcal{R} is compact, there is a global maximizer, by the Weierstrass theorem. Next lemma further simplifies the search by asserting that the boundary of \mathcal{R} contains a global maximizer. For the proof of Lemma 3, see [4].

Lemma 3 The boundary $\partial\mathcal{R}$ of the set \mathcal{R} contains a global maximizer of (10).

The boundary of \mathcal{R} is a closed curve in \mathbb{R}^2 . Our strategy consists in circulating along $\partial\mathcal{R}$ to spot a global maximizer. More precisely, we will sample $\partial\mathcal{R}$ with a finite set of points and pick the best point. To implement this strategy, we borrow the following theorem from [25].

Theorem 4 ([25]) Let $n \geq 3$ and let A, B be $n \times n$ symmetric matrices. Let

$$\mathcal{R}(A, B) = \left\{ (x, y) \in \mathbb{R}^2 : x = v^\top A v, y = v^\top B v, \text{ for some } v \in \mathbb{R}^n, v^\top v = 1 \right\}.$$

For $t \in [0, 2\pi]$, let $C(t) = A \cos t + B \sin t$ and let $\lambda_{\min}(t)$ be the minimal eigenvalue of the matrix $C(t)$ and $u_{\min}(t)$ an associated unit-norm eigenvector. Suppose that $\lambda_{\min}(t)$ is a simple eigenvalue of $C(t)$ for all $t \in [0, 2\pi]$. Then, the set $\mathcal{R}(A, B)$ is strictly convex and its boundary is given by

$$\partial\mathcal{R}(A, B) = \left\{ (x(t), y(t)) : t \in [0, 2\pi], x(t) = u_{\min}(t)^\top A u_{\min}(t), y(t) = u_{\min}(t)^\top B u_{\min}(t) \right\}.$$

Parametrization of $\partial\mathcal{R}$. Theorem 4 shows that the set $\mathcal{R}(A, B)$ (in our context, the set $\mathcal{R} = \mathcal{R}(S, mm^\top)$) can be parameterized by moving a single parameter t over the compact interval $[0, 2\pi]$. The theorem assumes that $\lambda_{\min}(t)$ is simple for all t , and, consequently, that $\mathcal{R}(A, B)$ is strictly convex. However, a parametrization of the boundary when this condition is not satisfied is readily available, as we explain next. Applied to our set \mathcal{R} in (12), this leads to the following procedure:

- 1) generate the points

$$(x_k, y_k) = (u_k^\top S u_k, u_k^\top m m^\top u_k), \quad k = 1, 2, \dots, K,$$

where u_k denotes an unit-norm eigenvector corresponding to the minimal eigenvalue of

$$C_k = S \cos((k-1)2\pi/K) + m m^\top \sin((k-1)2\pi/K).$$

Here, K is the user-defined grid size and $\{(x_k, y_k) : k = 1, \dots, K\}$ is an initial sample of $\partial\mathcal{R}$;

- 2) if the distance between two consecutive points (x_k, y_k) and (x_{k+1}, y_{k+1}) is greater than a prescribed threshold, interpolate the line segment which connects them, i.e., consider

$$\left(x_k^{(j)}, y_k^{(j)} \right) = (1 - j/J) (x_k, y_k) + j/J (x_{k+1}, y_{k+1}), \quad j = 0, 1, \dots, J.$$

In summary, our sampling of $\partial\mathcal{R}$ is $\widehat{\partial\mathcal{R}} = \{(x_k, y_k)\} \cup \{x_k^{(j)}, y_k^{(j)}\}$.

Solving (9). We explained how to solve (10) by parameterizing $\partial\mathcal{R}$. Now we explain how a solution of (10), x^*, y^* , is sufficient to reconstruct v^* , a solution of (9). Vector v^* in Lemma 2 can be found as follows. Let

$$(x^*, y^*) \in \arg \max_{(x, y) \in \widehat{\partial\mathcal{R}}} \psi_{\text{KL}}(x, y).$$

That is, (x^*, y^*) denotes the best point in $\widehat{\partial\mathcal{R}}$. If $(x^*, y^*) \in \{(x_k, y_k)\}$, say $(x^*, y^*) = (x_{k^*}, y_{k^*})$, then we can take v^* as an unit-norm eigenvector associated with the minimal eigenvalue of C_{k^*} . Otherwise, $(x^*, y^*) \in \{(x_k^{(j)}, y_k^{(j)})\}$, and we need to solve the system of 3 quadratic equations:

$$v^\top S v = x^*, \quad v^\top m m^\top v = y^*, \quad v^\top v = 1, \quad (13)$$

with respect to (w.r.t.) v . Any solution can be taken as v^* . It can be shown that (13) can be efficiently solved by solving a convex problem.

1.2) Case $p > 1$: greedy algorithm Optimization problem (8) for the case $p > 1$ is very difficult to solve globally; we propose a greedy, suboptimal approach. We construct the columns of the matrix $E = [e_1 \ e_2 \ \dots \ e_p]$ one by one (in the order e_1, e_2, \dots). We construct the j -th column by solving (9), with the constraint that the column e_j must be orthogonal to the previously determined columns e_1, e_2, \dots, e_{j-1} , i.e., we solve:

$$\begin{aligned} & \text{maximize} \quad \min_{m_0 \in \mathcal{E}(\widehat{m}_0, k_0 S_0^{-1}), m_1 \in \mathcal{E}(\widehat{m}_1, k_1 S_1^{-1})} f_{\text{KL}}(e, m_0, m_1, S_0, S_1) \\ & \text{subject to} \quad e^\top e = 1 \\ & \quad \quad \quad e^\top e_i = 0, \quad i = 1, \dots, j-1. \end{aligned} \quad (14)$$

Let $U^{(j)} \in \mathbb{R}^{n \times (n-j+1)}$ be a matrix with orthonormal columns which spans the orthogonal complement of $\text{span}\{e_1, \dots, e_{j-1}\}$. The restrictions in (14) mean that $e = U^{(j)} e^{(j)}$ for some unit-norm $e^{(j)} \in \mathbb{R}^{n-j+1}$.

This means that (14) corresponds to

$$\begin{aligned} & \text{maximize} \quad \min_{m_0 \in \mathcal{E}(\widehat{m}_0, k_0 S_0^{-1}), m_1 \in \mathcal{E}(\widehat{m}_1, k_1 S_1^{-1})} f_{\text{KL}}(U^{(j)} e^{(j)}, m_0, m_1, S_0, S_1) \\ & \text{subject to} \quad e^{(j)\top} e^{(j)} = 1. \end{aligned} \quad (15)$$

The problem (15) is equivalent to (16) (see [4])

$$\begin{aligned} & \text{maximize} \quad \min_{m_0^{(j)} \in \mathcal{E}(\widehat{m}_0^{(j)}, k_0 (S_0^{(j)})^{-1}), m_1^{(j)} \in \mathcal{E}(\widehat{m}_1^{(j)}, k_1 S_1^{(j)-1})} f_{\text{KL}}(e^{(j)}, m_0^{(j)}, m_1^{(j)}, S_0^{(j)}, S_1^{(j)}) \\ & \text{subject to} \quad e^{(j)\top} e^{(j)} = 1 \end{aligned} \quad (16)$$

where $\widehat{m}_i^{(j)} = U^{(j)\top} \widehat{m}_i$ and $S_i^{(j)} = U^{(j)\top} S_i U^{(j)}$, $i = 0, 1$. That is, (16) is simply an instance of (9)

in the reduced dimensional space \mathbb{R}^{n-j+1} , for which we have developed a global solution. Algorithm 1 outlines the overall approach.

Algorithm 1: Greedy algorithm

- 1: **for** $j = 1$ to p **do**
 - 2: Compute $U^{(j)} \in \mathbb{R}^{n \times (n-j+1)}$ ($U^{(1)} := I_n$), an orthonormal basis for the orthogonal complement of the $j - 1$ dimensional subspace $\text{span}\{e_1, e_2, \dots, e_{j-1}\}$
 - 3: Compute the projected means and covariances $m_i^{(j)} = U^{(j)\top} m_i$, $S_i^{(j)} = U^{(j)\top} S_i U^{(j)}$ for $i = 0, 1$
 - 4: Compute $S^{(j)} = (S_0^{(j)})^{-1/2} S_1^{(j)} (S_0^{(j)})^{-1/2}$, $m^{(j)} = (S_0^{(j)})^{-1/2} (m_1^{(j)} - m_0^{(j)})$, $M^{(j)} = m^{(j)} m^{(j)\top}$
 - 5: Solve (10) for $(M, S) := (M^{(j)}, S^{(j)})$; find $e^{(j)} \in \mathbb{R}^{n-j+1}$ as: $e^{(j)} := e^*$, as in Lemma 2
 - 6: Compute the j th column of E as $e_j = U^{(j)} e^{(j)}$
 - 7: **end for**
-

2) *Projection phase:* Relaxation phase III-A1 produces a Stiefel matrix E . Now, we project the matrix E back to the set of 0/1 selection matrices. We remark that the objective function $f_{\text{KL}}(\cdot)$ in eqn. (6) depends on the matrix E only through its range space; that is, $f_{\text{KL}}(EQ) = f_{\text{KL}}(E)$, for any Stiefel matrix E and for any orthogonal $p \times p$ matrix Q . Thus, we choose the selection matrix \tilde{E} with the range space closest to the range space of matrix E . It can be shown (see [3]) that \tilde{E} can be efficiently obtained as follows: if (j_1, j_2, \dots, j_p) denote the indices of the largest entries on the diagonal of EE^\top , then $\tilde{E} = [h_{j_1} h_{j_2} \dots h_{j_p}]$ where h_j stands for the j -th column of the identity matrix I_n . Thus, the Projection phase has very small computational cost.

3) *Refinement phase:* Once the projection to the set of 0/1 selection matrices is done and the matrix \tilde{E} is obtained, we finalize our algorithm with a local maximization around \tilde{E} to get E^* (see [14], [19] for very similar local searches.) Namely, for a given selection matrix E in the neighborhood of \tilde{E} , we find

$$f_{\text{KL, worst}}(E) := \min_{m_0 \in \mathcal{E}(\hat{m}_0, k_0 S_0^{-1}), m_1 \in \mathcal{E}(\hat{m}_1, k_1 S_1^{-1})} f_{\text{KL}}(E, m_0, m_1).$$

The procedure has p steps. We start with the matrix $E := \tilde{E}$. In the first step, all columns of the current selection matrix E are fixed except the first one, which is viewed as an optimization variable. The first column is swept through all canonical vectors h_j , $j = 1, \dots, n$, different from the remaining $p - 1$ columns of E . After all possible choices for the first column are tested, the column is frozen to the choice that gives the maximal $f_{\text{KL, worst}}$. In the second step, this procedure is repeated for the second column, and so on, up to the p -th step; after the p -th step is done, we set $E^* := E$.

We obtain the quantity $f_{\text{KL, worst}}(E)$ by first finding:

$$\min_{\substack{m_0 \in \mathcal{E}(\hat{m}_0, k_0, S_0^{-1}), \\ m_1 \in \mathcal{E}(\hat{m}_1, k_1, S_1^{-1})}} (m_1 - m_0)^\top E (E^\top S_0 E)^{-1} E^\top (m_1 - m_0), \quad (17)$$

and then adding the remaining terms of the function $f_{\text{KL}}(\cdot)$ that do not depend on m_1 and m_0 (see eqn. (6)).

Similarly as with the equivalence of (15) and (16), it can be shown that the minimum in eqn. (17) (n -dimensional problem) equals the following minimum (of corresponding p -dimensional problem):

$$\min_{m'_1 \in \mathcal{E}(Q^\top (E^\top S_0 E)^{-1/2} E^\top \hat{m}_0, k_0, I_p), m'_0 \in \mathcal{E}(Q^\top (E^\top S_0 E)^{-1/2} E^\top \hat{m}_1, k_1, \Lambda^{-1})} \|m'_1 - m'_0\|, \quad (18)$$

where Λ and Q are, respectively, the matrix of eigenvalues and the matrix of eigenvectors of $(E^\top S_0 E)^{-1/2} E^\top S_1 E (E^\top S_0 E)^{-1/2}$, and $\|\cdot\|$ denotes Euclidean norm. The problem (18) is a (convex) quadratically constrained quadratic problem (QCQP), and it can be solved with complexity $\mathcal{O}(p^3)$ (see, e.g., [26]).

B. Algorithm for the robust Chernoff-based selection: the R-C algorithm

In this subsection, we present the algorithm R-C, the Chernoff based sensor selection under the presence of uncertainties. As mentioned previously, we adopt the same methodology for solving both (5) and (4) and, consequently, the structure of R-C is the same as the one in R-KL. However, the problem (5) is more difficult than (4), due to the additional maximization over the parameter s . This will result in several specificities in R-C compared to R-KL. We present R-C by focusing on these specificities, phase by phase, whereas the overall structure remains the same as in R-KL.

The main difference between R-KL and R-C is in the Relaxation phase in the case $p = 1$. As with the KL case, we solve the resulting Chernoff problem globally; we next explain a solution.

1) *Relaxation phase: case $p = 1$: global solution:* Stiefel relaxation of (5) for $p = 1$ is given by:

$$\begin{aligned} & \text{maximize} \quad \min_{m_0 \in \mathcal{E}(\hat{m}_0, k_0, S_0^{-1}), m_1 \in \mathcal{E}(\hat{m}_1, k_1, S_1^{-1})} \max_{s \in [0,1]} f_C(s, e, m_0, m_1) \\ & \text{subject to} \quad e^\top e = 1, \end{aligned} \quad (19)$$

where

$$f_C(s, e, m_0, m_1) = \frac{s(1-s)}{2} \frac{(e^\top (m_1 - m_0))^2}{se^\top S_0 e + (1-s)e^\top S_1 e} - \frac{1}{2} \log \frac{(e^\top S_0 e)^s (e^\top S_1 e)^{1-s}}{se^\top S_0 e + (1-s)e^\top S_1 e}.$$

The first reformulation of (19) that we make is the conversion of the minimax into maximin problem, as Lemma 5 explains.

Lemma 5 Problem (19) is equivalent to:

$$\begin{aligned} & \text{maximize} \quad \max_{e^\top e=1} \min_{m_0 \in \mathcal{E}(\hat{m}_0, k_0 S_0^{-1}), m_1 \in \mathcal{E}(\hat{m}_1, k_1 S_1^{-1})} f_C(s, e, m_0, m_1) \\ & \text{subject to} \quad s \in [0, 1]. \end{aligned} \quad (20)$$

Proof: Function f_C is convex w.r.t. m_0 and m_1 , and concave w.r.t. s , and the constraint sets are compact and convex. Thus, the equivalence follows by Sion's minimax theorem ([27]). ■

Next, we focus on the inner maximization in (20):

$$\begin{aligned} & \text{maximize} \quad \min_{m_0 \in \mathcal{E}(\hat{m}_0, k_0 S_0^{-1}), m_1 \in \mathcal{E}(\hat{m}_1, k_1 S_1^{-1})} f_C(s, e, m_0, m_1) \\ & \text{subject to} \quad e^\top e = 1 \end{aligned}, \quad (21)$$

where $s \in [0, 1]$ is fixed. The following lemma is the counterpart of Lemma 2.

Lemma 6 Suppose (x^*, y^*) solves

$$\begin{aligned} & \text{maximize} \quad \psi_C(s, x, y) \\ & \text{subject to} \quad (x, y) \in \mathcal{R} \end{aligned} \quad (22)$$

where

$$\psi_C(s, x, y) = \frac{s(1-s)}{2} \frac{\left\{ \left(\sqrt{y} - \frac{1}{\sqrt{k_1}} \sqrt{x} - \frac{1}{\sqrt{k_0}} \right)^+ \right\}^2}{s + (1-s)x} - \frac{1}{2}(1-s) \log x + \frac{1}{2} \log(s + (1-s)x),$$

and $s \in [0, 1]$. Let $v^* \in \mathbb{R}^n$ be a unit norm vector that generates x^* and y^* , i.e., $x^* = v^{*\top} S v^*$ and $y^* = v^{*\top} m m^\top v^*$. Then, $e^* := S_0^{-1/2} v^* / \|S_0^{-1/2} v^*\|$ solves (21).

Proof: The proof is similar to the proof of Lemma 2 and is omitted. ■

By similar analysis as in subsection III-A1, it can be shown that (22) can be solved by parametrization of the boundary of \mathcal{R} , given in eqn. (12). Thus, for fixed s , the algorithm that solves (21) is the same as the one that solves (9), except that $\psi_{KL}(x, y)$ is replaced by $\psi_C(s, x, y)$. Then, the function $\psi_C(s, x^*(s), y^*(s))$ can be evaluated using this algorithm and problem (19) is solvable by, e.g., a grid search on the interval $[0, 1]$.

The projection phase of R-C is the same as the projection phase of R-KL and the steps in the refinement phase of R-C are the same as the ones in the refinement phase of R-KL (with $f_{KL}(E, m_0, m_1)$ replaced by $\max_{s \in [0, 1]} f_C(s, E, m_0, m_1)$). Similarly as in III-A3, in the refinement phase, for a given selection matrix E , we have to find

$$f_{C, \text{worst}}(E) := \min_{m_0 \in \mathcal{E}(\hat{m}_0, k_0 S_0^{-1}), m_1 \in \mathcal{E}(\hat{m}_1, k_1 S_1^{-1})} \max_{s \in [0, 1]} f_C(s, E, m_0, m_1).$$

Applying again the minimax theorem, we first exchange the order of min and max in $f_{C, \text{worst}}$; then, for fixed s , we find

$$\min_{\substack{m_0 \in \mathcal{E}(\hat{m}_0, k_0 S_0^{-1}), \\ m_1 \in \mathcal{E}(\hat{m}_1, k_1 S_1^{-1})}} (m_1 - m_0)^\top E \left(s E^\top S_0 E + (1 - s) E^\top S_1 E \right)^{-1} E^\top (m_1 - m_0), \quad (23)$$

by solving the equivalent p -dimensional QCQP (24) (with equal minimum):

$$\min_{\substack{m'_0 \in \mathcal{E}(Q^\top (E^\top S_0 E)^{-1/2} E^\top \hat{m}_0, k_0 I_p), \\ m'_1 \in \mathcal{E}(Q^\top (E^\top S_1 E)^{-1/2} E^\top \hat{m}_1, k_1 \Lambda^{-1})}} (m'_1 - m'_0)^\top (s I_p + (1 - s) \Lambda)^{-1} (m'_1 - m'_0). \quad (24)$$

Finally, we compute $f_{C, \text{worst}}(E)$ by the bisection method on parameter s .

C. Complexities of R-KL and R-C

The complexity of both R-KL and R-C algorithm is $\mathcal{O}(n^3 p + np^4)$, although the hidden constant in R-C is larger than the one in R-KL. The least computational effort is required for the projection phase 2, which for both R-C and R-KL is $\mathcal{O}(n^2)$ and is dominated by complexities of the other two phases. It can be shown that Phase 1 has complexity $\mathcal{O}(n^3 p)$ and Phase 3 has complexity $\mathcal{O}(np^4)$.

IV. SENSOR SELECTION ALGORITHMS: NO UNCERTAINTIES CASE

In this section, we address a special, yet important case, when there are no uncertainties in the mean vectors and the problems (4) and (5) simplify by dropping the inner minimizations. We first remark that algorithms R-KL and R-C can readily solve the simplified versions of (4) and (5). However, we derive in this section a more efficient algorithm. We exploit the structure of the problem and the knowledge of exact distribution parameters m_i , $i = 0, 1$ (more specifically, their difference $m_1 - m_0$) to reduce the computational load of the relaxation phase of R-KL and R-C algorithm, while keeping the second and the third phase the same². The key to reducing the complexity of the relaxation phase is a *simple, analytic solution* of the relaxed, Stiefel problem, in the special case of equal mean values. We refer to the overall simplified algorithms as MD-KL (mean-difference based KL algorithm), and MD-C (mean-difference based C algorithm).

A. Kullback-Leibler based selection without uncertainties: The mean-difference KL algorithm (MD-KL)

In this subsection, we only explain the relaxation phase of MD-KL, as the other two phases are the same as in R-KL. We first consider a special case of equal mean values of the problem (8) (without

²We remark that the refinement phase of R-C and R-KL simplifies significantly in the no uncertainties case, as in this case computing $f_{\text{KL}, \text{worst}}(E)$ boils down to computing $f_{\text{KL}}(E)$, i.e. there is no need to solve intermediate minimization problems (see (17) and (23)).

inner minimization), and we show that this problem has a simple analytic solution. Based on the solution for the equal means, we derive an algorithm that solves the general case.

1) *Relaxation phase: the case $m_0 = m_1$:* Consider the problem (8) when there is no uncertainty in the mean values (uncertainty ellipsoids shrink to a point, by letting $k_0, k_1 = \infty$, and inner minimization drops from the problem). We first remark that we can write (8) when $k_i = \infty, m_1 = m_0$, in the following form:

$$\begin{aligned} & \text{maximize} && \frac{1}{2} (\text{tr}(P^\top S P) - \log |P^\top S P| - p) \\ & \text{subject to} && P^\top P = I_p, \end{aligned} \quad (25)$$

where $S = S_0^{-1/2} S_1 S_0^{-1/2}$. This equivalence can be shown by noting that the constraint $E^\top E = I_p$ in (8) can be replaced by $E^\top S_0 E = I_p$, and by introducing the new variable $P = S_0^{1/2} E$.

The objective function in (25) can be further simplified to $\sum_{i=1}^p \phi_{\text{KL}}(\lambda_i(P^\top S P))$, where $\phi_{\text{KL}}(x) = x - \log x - 1$ and λ_i denotes the i -th largest eigenvalue. Invoking Poincaré separation theorem (see [28]), we show in the Appendix C that the solution P^* of (25) is given by a set of p orthonormal eigenvectors of S corresponding to the p eigenvalues that give the highest objective $\phi_{\text{KL}}(\cdot)$. It turns out that we do not need to check all $\binom{n}{p}$ eigenvector subsets, but only at most $p+1$ of them. We now give the procedure in Algorithm 2 to choose the optimal subset of eigenvectors (see Appendix C for the proof).

Algorithm 2: Procedure for solving (25) when $m_1 = m_0$

- 1: Set $\phi_{\text{KL}}^* = 0$
 - 2: **for** $j = 0$ to p **do**
 - 3: $x = (\lambda_1(S), \dots, \lambda_j(S), \lambda_{n-p+j+1}(S), \dots, \lambda_n(S))$
 - 4: Compute $\phi = \sum_{i=1}^p \phi_{\text{KL}}(x_i)$
 - 5: if $\phi > \phi_{\text{KL}}^*$, then $j^* = j, \phi_{\text{KL}}^* = \phi$
 - 6: **end for**
 - 7: $x^* = (\lambda_1(S), \dots, \lambda_{j^*}(S), \lambda_{n-p+j^*+1}(S), \dots, \lambda_n(S))$;
 P^* is the set of eigenvectors of S corresponding to the eigenvalues of S given in x^* .
-

We give the intuition behind the solution of (25). Recall that the matrix E is chosen such that the projection of the covariance matrix S_0 equals $E^\top S_0 E = I_p$. Then, the projection of S_1 , $E^\top S_1 E$, equals $P^\top S P$. Thus, the further from point 1 are the eigenvalues of $P^\top S P$, the better are the projected distributions separated. The function $\phi_{\text{KL}}(\cdot)$ measures the distance from 1 in this sense.

2) *Relaxation phase: general case:* The main idea behind the relaxation phase of MD-KL is as follows: set one column of the solution Stiefel matrix E in the direction of the vector $m_1 - m_0$, i.e., in the direction of the difference of the distribution means. The remaining $p - 1$ columns of E are then obtained in the following way: we project the distribution parameters $m_i, S_i, i = 0, 1$, to the orthogonal

complement of $m_1 - m_0$, and then solve an $(p - 1)$ dimensional instance of (8) when $k_i = \infty$, with the projected distribution parameters. This $p - 1$ -dimensional instance of (8) is in fact the special one, with equal means, and is hence very efficiently solved by procedure given in Algorithm (2). The relaxation phase of MD-KL is summarized in Algorithm 3. We give the intuition behind the choice of $m_1 - m_0$ as the direction of the first column of E : the Euclidian distance between the means $E^\top m_1$ and $E^\top m_0$ of the projected distributions is maximal possible (and equal to $\|m_1 - m_0\|$), when one column of E lies in the direction of $m_1 - m_0$.

Algorithm 3: MD-KL algorithm: Relaxation phase

- 1: Set $e_1 := (m_1 - m_0) / \|m_1 - m_0\|$
 - 2: Compute $U \in \mathbb{R}^{n \times (n-1)}$, an orthonormal basis for the orthogonal complement of e_1
 - 3: Compute $S' = (U^\top S_0 U)^{-1/2} U^\top S_1 U (U^\top S_0 U)^{-1/2}$
 - 4: Find P as in Algorithm 2 for $S := S'$, and $p := p - 1$
 - 5: $[e_2, e_3, \dots, e_p] := U S_0^{-1/2} P$, $E := [e_1, e_2, \dots, e_p]$
-

B. Chernoff based selection without uncertainties: The mean-difference C algorithm (MD-C)

1) *Relaxation phase: the case $m_0 = m_1$:* The counterpart of problem (25) for the C criterion is the following:

$$\begin{aligned} & \text{maximize} \quad \max_{s \in [0,1]} -\log \frac{|P^\top S P|^{1-s}}{|sI + (1-s)P^\top S P|}, \\ & \text{subject to} \quad P^\top P = I_p \end{aligned} \quad (26)$$

The objective in (26) can be written as $\sum_{i=1}^p \phi_C(s, \lambda_i(P^\top S P))$, where $\phi_C(s, x) := \log(s + (1-s)x) - (1-s) \log x$. Now, the problem (26) ($m_1 = m_0$) can be solved by a method similar to the one in IV-A1. Namely, 1) by invoking Poincaré separation theorem; 2) by using concavity of $\phi_C(\cdot, x)$; and 3) by using unimodularity of $\phi_C(s, \cdot)$, it can be shown that the solution to (26) ($m_1 = m_0$) can be obtained by the procedure given in Algorithm 4.

General case of the relaxation phase of MD-C algorithm is the same as the one in MD-KL given in 3, except that, in step 4), it calls procedure given by Algorithm 4.

C. Complexities of MD-KL and MD-C

We briefly comment on the complexity of MD-KL and MD-C. The complexity of both MD-KL and MD-C is $\mathcal{O}(n^3 + np^3)$, although the hidden constant in MD-C is larger than the one in MD-KL. The main computational burden in the relaxation phase of MD-KL is computation of the orthogonal complement U of $m_1 - m_0$, and subsequent eigenvalue decomposition of the supplementary matrix S (see Algorithm 3),

Algorithm 4: Procedure for solving (26) when $m_1 = m_0$

- 1: Set $\phi_C^* = 0$
 - 2: **for** $j = 0$ to p **do**
 - 3: $x = (\lambda_1(S), \dots, \lambda_j(S), \lambda_{n-p+j+1}(S), \dots, \lambda_n(S))$
 - 4: Find $s(x) \in \arg \max_{s \in [0,1]} \sum_{i=1}^p \phi_C(s, x_i)$ by Newton method
 - 5: Compute $\phi = \sum_{i=1}^p \phi_C(s(x), x_i)$
 - 6: if $\phi > \phi_C^*$, then $j^* = j$, $s^* = s(x)$, $\phi_C^* = \phi$
 - 7: **end for**
 - 8: $x^* = (\lambda_1(S), \dots, \lambda_{j^*}(S), \lambda_{n-p+j^*+1}(S), \dots, \lambda_n(S))$; $s^* = s^*$;
 P^* is the set of orthonormal eigenvectors of S corresponding to the eigenvalues of S given in x^* .
-

which is in total of order $\mathcal{O}(n^3)$. The relaxation phase of MD-C requires more computational effort; besides finding U and the eigenvalue decomposition of S , each of the $p - 1$ steps in Algorithm 4 requires finding optimal s by Newton method (contrary to just evaluating the KL distance for a given choice of $p - 1$ eigenvalues of S with Algorithm 2). However, the number of operations in $p - 1$ Newton runs is still dominated by the number of operations to find U . Therefore, the complexity of the relaxation phase of MD-C is $\mathcal{O}(n^3)$. Finally, it can be shown that, for both MD-KL and MD-C, the refinement phase is of complexity $\mathcal{O}(np^3)$.

V. NUMERICAL STUDIES: TESTING THE OPTIMIZATION CRITERIA

This subsection tests how good are the Kullback-Leibler and the Chernoff distance as optimality criteria for sensor selection. To this end, we want to compare the sensor selections that optimize the KL and C distances with: 1) the sensor selection that minimizes the Bayes probability of error (Bayes optimality); 2) the sensor selection that minimizes the probability of miss subject to a given probability of false alarm (Neyman-Pearson optimality). We find numerically the Bayes optimal and the Neyman-Pearson optimal sensor selections by Monte Carlo simulations. With respect to Bayes optimality, for each possible (out of $\binom{n}{p}$) sensor selection, we estimate the Bayes probability of error P_e by Monte Carlo simulations with 100,000 instantiations of the maximum-likelihood detector tests (with zero threshold), with equal prior probabilities. With respect to Neyman-Pearson optimality, for the fixed probability of false alarm P_{FA} , we find the sensor selection that maximizes the probability of detection P_D . This is achieved by estimating the receiver operating characteristic (ROC) curve (by Monte Carlo simulations with 20,000 instantiations of the likelihood ratio tests) in the neighborhood of P_{FA} in a fine grid, and then interpolating the ROC curve (i.e., P_D) at the desired point P_{FA} . This is done for each possible selection and the selection with maximal obtained P_D is set as Neyman-Pearson optimal.

Table I shows the Bayes probability of error for: 1) the sensor selection that maximizes KL distance (KL); 2) the sensor selection that maximizes C distance (C); 3) the best sensor selection (that minimizes P_e); 4) the worst sensor selection (that maximizes P_e); and 5) the average P_e over all selections. Table I (left) is for $n = 12$ and $p = 2, 3, 4, 5$ and Table I (right) is for $n = 15$ and $p = 2, 3, 4, 6$. We can see that C selection matches the best selection. KL selection is in many cases very close or equal to C and best selections in P_e . As could be predicted by theory (Chernoff lemma), the C selection is better than the KL selection in terms of P_e .

Comparison of: 1) C, KL and best selections; with 2) worst and average selections justifies the sensor selection problem; namely, by finding the optimal selection, P_e can be for an order of magnitude smaller than for the average selection. (See Table 1 and Figure 1-right.)

Table II shows the probability of detection for: 1) the KL-optimal selection (KL); 2) the C-optimal selection (C); and 3) the Neyman-Pearson (NP) optimal selection, for $n = 12$, $p = 2, 3, 4, 5$. We can see that, for $p = 2$ and $p = 3$, both KL and C selections match the NP optimal selections, and for $p = 3$ and $p = 4$, P_D for KL and C selections is at most 3.5% from the optimum ($p = 4$, $P_D = 0.05$).

Figure 1 (left) plots the ROC curves for all possible selections, for $n = 5$ and $p = 2$. We plot the ROC curves for: 1) KL-optimal selection; 2) C-optimal selection; 3) the pointwise envelope of all possible curves (Neyman-Pearson optimal). Remark that the Neyman-Pearson optimal curve (envelope) is not obtained for a single selection; in different regions, it corresponds to different selections. We can see that for lower values of P_{FA} , KL selection is optimal; for higher values of P_D , C-selection is optimal.

Figure 1 plots the ROC curve for all possible $\binom{n}{p} = 1365$ selections for a larger example, with $n = 15$ and $p = 4$. Interestingly, we can see that the KL and C selections are very close to the optimum, in whole range of P_{FA} . In addition, we plot the average of the ROC curves (pointwise average of P_D for each fixed P_{FA}). This average curve thus represents what performance would be, on average, achieved, if we choose a subset of sensors uniformly at random. We can see that there is a large gain of the C and KL selections over this average curve; thus, selecting the optimal, rather than random subset of sensors, provides large performance gain.

Finally, we remark that, in extensive simulations, we observe similar behavior as in representative Tables I and II, and Figure 1 (left and right). That is, the KL and C selections are very close to optimal and even equal to optimal in certain range of P_{FA} . We also report that C-selection is generally better than KL for large P_D 's (upper right part of the ROC curve,) while KL is generally better for low P_{FA} (lower left in the ROC). Improvement of KL over C for low P_{FA} is smaller than the improvement of C over KL for large P_D .




Fig. 1. Left: $n = 5$, $p = 2$; ROC curves for all possible sensor selections (gray); the ROC curves for the KL-optimal and the C-optimal sensor selections, as well as the Neyman-Pearson optimal “envelope” are represented in different color. Right: $n = 15$, $p = 4$, Figure plots the same ROC curves as on the left; in addition, the Figure plots the average ROC among all sensor selections.

TABLE I
 P_e FOR KL, C AND OPTIMAL SELECTION; LEFT: $n = 12$, $p = 2, 3, 4, 5$; RIGHT: $n = 15$, $p = 2, 3, 4, 6$

P_e	$p = 2$	$p = 3$	$p = 4$	$p = 5$
KL	0.100	0.078	0.052	0.046
C	0.100	0.078	0.052	0.043
Bayes-best	0.100	0.078	0.052	0.043
worst	0.457	0.439	0.396	0.319
average	0.275	0.216	0.170	0.134

P_e	$p = 2$	$p = 3$	$p = 4$	$p = 6$
KL	0.086	0.061	0.037	0.022
C	0.085	0.051	0.029	0.012
Bayes-best	0.085	0.051	0.029	0.012
worst	0.480	0.440	0.396	0.311
average	0.240	0.180	0.136	0.077

VI. NUMERICAL STUDIES: TESTING THE ALGORITHMS

This subsection tests our algorithms for solving the sensor selection problem. Subsection VI-A shows that algorithms R-KL and R-C show near optimal performance in the presence of uncertainties (i.e., for solving the problems (4) and (5)), Subsection (VI-B) shows that R-KL and R-C have near optimal

TABLE II
 P_D FOR $P_{FA} = [0.005 \ 0.03 \ 0.1]$, $n = 15$, $p = 2, 3, 4, 5$ FOR KL, C AND OPTIMAL SELECTION

P_D	$P_{FA} = 0.005$			$P_{FA} = 0.03$			$P_{FA} = 0.1$		
	KL	C	NP	KL	C	NP	KL	C	NP
$p = 2$	0.720	0.720	0.720	0.889	0.889	0.889	0.961	0.961	0.961
$p = 3$	0.812	0.812	0.812	0.941	0.941	0.941	0.985	0.985	0.985
$p = 4$	0.868	0.838	0.868	0.962	0.967	0.970	0.990	0.995	0.995
$p = 5$	0.891	0.898	0.903	0.968	0.981	0.983	0.991	0.994	0.996

performance when also applied to the case of no uncertainties (i.e., when $k_0 = k_1 = \infty$ in (4) and (5)). This subsection also shows that MD-KL (resp. MD-C) has comparable performance to R-KL (resp. R-C), while reducing the computational time.

For smaller problem instances, i.e., for smaller values of n and p , we compare the solutions produced by our algorithms with the optimum obtained by exhaustive search. For larger examples, when the optimum is infeasible to compute, we generate randomly a number of sensor selections; then, we compare the selection obtained by our algorithms with the best among generated random selections.

For a fixed size of the problem (for fixed pair (n, p)), we generate randomly K_{exper} instances of the problem parameters $(m_i, S_i, i = 0, 1)$; for the case with uncertainties, $K_{\text{exper}} = 50$, and for the case without uncertainties, $K_{\text{exper}} = 200$. In the case with uncertainties, parameters k_0 and k_1 are chosen such that the norm of the mean vector drift does not exceed $0.15\|m_1 - m_0\|$, i.e., we set them as $k_i = |S_i| / (0.15\|m_1 - m_0\|)^2$, $i = 0, 1$. Based on K_{exper} solved problems of fixed size (n, p) , we create statistics on how our algorithms behave for given size of the problem.

A. Testing the R-KL and R-C algorithms: The case with uncertainties

Testing the robustness against the uncertainty in distribution means. Table III (left) shows performance of the R-KL and R-C algorithms, for solving (4) and (5), respectively. We evaluate the optimal solutions $f^{\text{OPT-KL}}$ and $f^{\text{OPT-C}}$ by exhaustive search. We then compute the ratio $r_{\text{R-KL}} := \frac{f^{\text{R-KL}}}{f^{\text{OPT-KL}}}$ (resp. $r_{\text{R-C}} := \frac{f^{\text{R-C}}}{f^{\text{OPT-C}}}$) that says how close is the solution value obtained by R-KL (resp. R-C) to optimum. Table III (left) shows the results for $n = 10, 12, 15$ and $p = 3$. We present the maximum (max), the average (avg), the minimum (min), and the standard deviation (dev) of the ratio $r_{\text{R-KL}}$ over $K_{\text{exper}} = 50$ problem instances, for each pair (n, p) . We can see that both R-KL and R-C show very good performance; the R-C shows better performance than R-KL. The maximal value of both $r_{\text{R-KL}}$ and $r_{\text{R-C}}$ in all the examples is equal to 1. With the R-KL algorithm, the average value of $r_{\text{R-KL}}$ is always above 91.8%, and the minimum value is always above 51%; with the R-C algorithm, the average of $r_{\text{R-C}}$ is always above 96%, and the minimum is always above 59%.

Table III (right) shows performance of R-KL and R-C for a larger example $n = 50$, $p = 5$, for which it was infeasible to find the optimal solution by exhaustive search. We thus randomly generate 2500 sensor selections, and we evaluate the quantity $f^{\text{BEST-RAND-KL}}$ —the maximal objective function in (4) over all 2500 randomly generated selections. Define the ratios $\rho_{\text{R-KL}} := \frac{f^{\text{R-KL}}}{f^{\text{BEST-RAND-KL}}}$. (Analogously define the corresponding ratios for the Chernoff-based sensor selection.) We report that, to obtain $f^{\text{BEST-RAND-KL}}$ (for 2500 50×5 selection matrices) it takes about 10 times longer than for R-KL algorithm to find a

TABLE III

LEFT: STATISTICS FOR r_{R-KL} AND r_{R-C} , SMALLER EXAMPLES: $n = 10, 12, 15$, $p = 3$, $K_{\text{exper}} = 50$; RIGHT: STATISTICS FOR ρ_{R-KL} AND ρ_{R-C} , $n = 50$, $p = 5$, $K_{\text{exper}} = 50$

	max	avg	min	dev		max	avg	min	dev
r_{R-KL}									
$n = 10$	1.000	0.964	0.606	0.090					
$n = 12$	1.000	0.918	0.551	0.140					
$n = 15$	1.000	0.939	0.512	0.118					
r_{R-C}					ρ_{R-KL}	1.777	1.267	0.817	0.191
$n = 10$	1.000	0.981	0.7862	0.051	ρ_{R-C}	1.868	1.277	1.005	0.166
$n = 12$	1.000	0.982	0.834	0.037					
$n = 15$	1.000	0.961	0.595	0.095					

solution. Also, to obtain $f^{\text{BEST-RAND-C}}$ it takes about 6 times longer than for R-C algorithm to find a solution. From Table III (right) we can see that R-KL in many cases outperforms random strategy: with significant savings in time (90%) it gives a 26% larger objective function on average. R-C always stays above the random strategy in terms of the objective function, with computational savings of 83%.

B. Testing the algorithms: The case without uncertainties

Smaller examples: Comparison with the global optimum. Table IV shows performance of the R-KL and MD-KL algorithms for solving (4) in the case without uncertainties ($k_0 = k_1 = \infty$); table V shows performance of the R-C and MD-C algorithms for solving the corresponding Chernoff problem (4). For fixed n and p , we generate $K_{\text{exper}} = 200$ sets of problem parameters $m_i, S_i, i = 0, 1$; for each fixed $m_i, S_i, i = 0, 1$, we evaluate the ratios r_{R-KL} , r_{MD-KL} , r_{R-C} , and r_{MD-C} . Table IV shows the maximum (max), the average (avg) and the minimum (min) of the quantities r_{R-KL} and r_{MD-KL} over all $K_{\text{exper}} = 200$ experiments. The standard deviation in these experiments varies from 0.03 to 0.06. We can see from the table that the maximum of both r_{R-KL} and r_{MD-KL} is always 1. Table IV demonstrates very good performance of both R-KL and MD-KL algorithms. On average, r_{R-KL} and r_{MD-KL} are always (for each pair n, p) above 97.7%; the minimum is always above 67.2%. We see that R-KL and MD-KL algorithms have comparable performance with respect to (near)optimality, while MD-KL has smaller computational cost (We note, however, that MD-KL does not apply for the case with uncertainties in distribution means.)

Table V shows very good performance of the R-C algorithm and the MD-C algorithm. As with the

TABLE IV
STATISTICS FOR $r_{\text{R-KL}} = f^{\text{R-KL}}/f^{\text{OPT-KL}}$ AND $r_{\text{MD-KL}} = f^{\text{MD-KL}}/f^{\text{OPT-KL}}$

	$p = 3$			$p = 4$			$p = 5$		
	max	avg	min	max	avg	min	max	avg	min
$r_{\text{R-KL}}$									
$n = 20$	1.000	0.990	0.789	1.000	0.985	0.688	1.000	0.977	0.672
$n = 30$	1.000	0.989	0.830	1.000	0.988	0.826	1.000	0.983	0.795
$n = 40$	1.000	0.989	0.729	1.000	0.985	0.842	1.000	0.983	0.817
$r_{\text{MD-KL}}$									
$n = 20$	1.000	0.992	0.744	1.000	0.982	0.688	1.000	0.975	0.672
$n = 30$	1.000	0.989	0.809	1.000	0.987	0.832	1.000	0.981	0.742
$n = 40$	1.000	0.985	0.729	1.000	0.980	0.802	1.000	0.981	0.834

TABLE V
STATISTICS FOR $r_{\text{R-C}} = f^{\text{R-C}}/f^{\text{OPT-C}}$ AND $r_{\text{MD-C}} = f^{\text{MD-C}}/f^{\text{OPT-C}}$

	$p = 3$			$p = 4$			$p = 5$		
	max	avg	min	max	avg	min	max	avg	min
$r_{\text{R-C}}$									
$n = 20$	1.000	0.994	0.831	1.000	0.992	0.789	1.000	0.995	0.850
$n = 30$	1.000	0.997	0.880	1.000	0.997	0.868	1.000	0.995	0.883
$n = 40$	1.000	0.998	0.959	1.000	0.995	0.936	1.000	0.997	0.959
$r_{\text{MD-C}}$									
$n = 20$	1.000	0.997	0.835	1.000	0.995	0.874	1.000	0.996	0.918
$n = 30$	1.000	0.995	0.874	1.000	0.997	0.892	1.000	0.995	0.928
$n = 40$	1.000	0.998	0.931	1.000	0.994	0.933	1.000	0.994	0.953

case of KL algorithms, the maximum is always 1. The average of both quantities $r_{\text{R-C}}$ and $r_{\text{MD-C}}$ is always above 99.2%, and the minimum is always above 78.9%. Standard deviation is smaller than the one in R-KL and varies between 0.01 and 0.026. Thus, R-C and MD-C show closer near optimality than R-KL and MD-KL.

Larger examples: Comparison with the best randomly generated selection. We now consider larger n and p , when computing the optimum by the exhaustive search is infeasible. Similarly as in subsection VI-A (larger examples), we randomly generate 10^5 sensor selections and find the maximal objective function $f^{\text{BEST-RAND-KL}}$ over these selections. For a fixed pair n and p , we generate $K_{\text{exper}} = 100$ sets of the distribution parameters $m_i, S_i, i = 0, 1$; for each set of the distribution parameters, we evaluate $\rho_{\text{R-KL}}$,

TABLE VI
STATISTICS FOR $\rho_{R-KL} = f^{R-KL} / f^{BEST-RAND-KL}$ AND $\rho_{MD-KL} = f^{MD-KL} / f^{BEST-RAND-KL}$

	$p = 0.1 \times n$			$p = 0.2 \times n$			$p = 0.3 \times n$		
	avg	max	min	avg	max	min	avg	max	min
ρ_{R-KL}									
$n = 50$	1.069	1.246	0.855	1.225	1.577	1.072	1.274	1.733	1.071
$n = 80$	1.298	1.544	1.162	1.474	1.860	1.233	1.472	1.962	1.229
$n = 100$	1.427	1.705	1.169	1.556	1.739	1.337	1.598	1.894	1.378
ρ_{MD-KL}									
$n = 50$	1.072	1.246	0.919	1.223	1.464	1.002	1.265	1.684	1.077
$n = 80$	1.299	1.544	1.130	1.475	1.782	1.248	1.468	2.002	1.225
$n = 100$	1.429	1.705	1.169	1.563	1.864	1.384	1.617	1.973	1.374

ρ_{MD-KL} , and ρ_{R-C} , ρ_{MD-C} . We are also interested in the ratio of the computational time of the R-KL (resp. MD-KL) algorithm over the computational time of checking 10^5 random selections (and the analogous quantities for the Chernoff-based selection).

Table VI shows the average (avg), the maximal (max), and the minimal (min) values of ρ_{R-KL} and ρ_{MD-KL} over 100 data sets m_i , S_i , $i = 0, 1$; Table VII (left) shows the computational time ratios. We can see that both R-KL and MD-KL outperform best random selection strategy, as they achieve larger objective function while reducing computational time. For example, for $n = 80$, $p = 0.1 \times n = 8$, R-KL can give 54% larger objective function, while it has 3 times smaller computational time. MD-KL is better than R-KL, as it shows comparable performance in terms of the objective function; at the same time, it significantly reduces the computational time. Thus, for sensor selection without uncertainties, MD-KL is a very good tool that can handle large problem instances. Similar conclusions hold for the Chernoff based counterpart algorithms (see Table VII and Table VIII (right)). Chernoff based algorithms have larger computational time than Kullback-Leibler based algorithms, which is expected due to additional maximization over variable s (see (5)). Section V shows generally that Chernoff criterion has some advantage over Kullback-Leibler criterion, which trades off the computational requirements to solve for these criteria.

VII. CONCLUSIONS

In this paper, we addressed the problem of finding the most informative subset of p out of n sensors, for the task of deciding between the two possible hypothesis on the monitored environment. We proposed

TABLE VII
STATISTICS FOR $\rho_{R-C} = f^{R-C}/f^{\text{BEST-RAND-C}}$ AND $\rho_{MD-C} = f^{MD-C}/f^{\text{BEST-RAND-C}}$

	$p = 0.1 \times n$			$p = 0.2 \times n$			$p = 0.3 \times n$		
	avg	max	min	avg	max	min	avg	max	min
ρ_{R-C}									
$n = 50$	1.074	1.214	0.992	1.182	1.341	1.072	1.194	1.307	1.095
$n = 80$	1.262	1.454	1.132	1.357	1.621	1.199	1.338	1.464	1.206
$n = 100$	1.375	1.517	1.251	1.447	1.616	1.299	1.402	1.629	1.287
ρ_{MD-C}									
$n = 50$	1.074	1.214	0.991	1.182	1.341	1.038	1.195	1.307	1.100
$n = 80$	1.262	1.471	1.132	1.357	1.621	1.194	1.338	1.464	1.212
$n = 100$	1.375	1.517	1.254	1.445	1.616	1.296	1.403	1.632	1.296

TABLE VIII
AVERAGE TIME RATIOS, $K_{\text{exper}}=100$; LEFT: KL ALGORITHMS; RIGHT: C ALGORITHMS

	$p = 10\%$	$p = 20\%$	$p = 30\%$		$p = 10\%$	$p = 20\%$	$p = 30\%$
R-KL				R-C			
$n = 50$	0.070	0.104	0.107	$n = 50$	0.105	0.146	0.153
$n = 80$	0.276	0.295	0.291	$n = 80$	0.472	0.549	0.507
$n = 100$	0.417	0.472	0.349	$n = 100$	0.875	0.990	0.721
MD-KL				MD-C			
$n = 50$	0.003	0.005	0.006	$n = 50$	0.002	0.004	0.006
$n = 80$	0.007	0.012	0.014	$n = 80$	0.006	0.011	0.014
$n = 100$	0.011	0.017	0.021	$n = 100$	0.010	0.017	0.021

two different information theoretic criteria for the best sensor selection: the Kullback-Leibler distance and the Chernoff distance between the distributions induced by selected sensors. We tackled the case where the distributions are Gaussian, but the mean vectors are known only up to confidence regions. We formulated the corresponding maxmin optimization problems, and developed the R-KL and R-C algorithms, that efficiently solve the problems with complexity $\mathcal{O}(n^3p + np^4)$. We also addressed the case when the mean vectors are known and, for this case, we exploited the structure of the problems to develop more efficient algorithms, MD-KL and MD-C, of complexity $\mathcal{O}(n^3 + np^3)$. We performed Monte-Carlo based experiments to test both the proposed sensor selection criteria and the sensor selection

algorithms. Numerical studies of the criteria show that the Kullback-Leibler based and the Chernoff based selections have near optimal performance, both in the Neyman-Pearson and Bayes sense. The performance of our algorithms we compared with 1) the optimal sensor selections, when the exhaustive search to compute them is feasible (smaller n and p); and 2) with best random selections, when n and p are large. Comparison with the exhaustive search shows that proposed algorithms in many cases find the optimal selection and, on average, are at most 5% below the optimal value; at the same time, computational savings are significant. For larger problems, simulation results demonstrate that our algorithm outperforms random searches, once an upper bound on computational time is set.

APPENDIX

A. Proof of NP hardness of optimization problems (4) and (5)

We will prove that both problems (4) and (5) are NP hard by reducing them to the maximal clique problem (MQP), which is known to be NP hard, see, e.g., [29]. We first define the MQP. Consider an undirected, simple (i.e., without self-loops) graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where \mathcal{N} is the set of vertices with cardinality $|\mathcal{N}| = n$, and \mathcal{E} is the set of undirected edges $\{i, j\}$, $|\mathcal{E}| = m$. A clique of the graph \mathcal{G} , of size p , is a complete subgraph of \mathcal{G} that has p vertices. The decision version of the maximum clique problem is as follows:

MQP: The maximal clique problem: “For given graph \mathcal{G} and a positive integer p , $1 \leq p \leq n$, determine whether \mathcal{G} has a clique of size at least p .”

We conduct a reduction to MQP by attaching to a graph \mathcal{G} a $n \times n$ matrix $S(\mathcal{G})$ of a special structure. Namely, we define a positive definite matrix $S(\mathcal{G})$ as follows:

$$[S(\mathcal{G})]_{ij} = \begin{cases} 2n & ; \text{if } i = j \\ -1 & ; \text{if } i \neq j, \{i, j\} \in \mathcal{E} \\ 0 & ; \text{otherwise} \end{cases} \quad (27)$$

The matrix $S(\mathcal{G})$ is positive definite because it has positive diagonal elements and it is strictly diagonally dominant. Now, fix an integer p , $1 \leq p \leq n$, and consider a set of matrices \mathcal{A}_p defined as $\mathcal{A}_p = \{A \in \mathbb{R}^{p \times p} : A = A^\top, A_{ii} = 2n, \forall i, A_{ij} \in \{0, -1\}, i \neq j\}$. Clearly, all matrices in \mathcal{A}_p are positive definite, as they are strictly diagonally dominant, with positive diagonal entries. Denote with $\mathbf{1}_p$ the column vector with all entries equal to 1 and define the function **s.e.i.** : $\mathcal{A}_p \rightarrow \mathbb{R}$ as **s.e.i.**(A) = $\mathbf{1}_p^\top A^{-1} \mathbf{1}_p$ (**s.e.i.** is the sum of the elements in the inverse.) We have the following result on the matrices in \mathcal{A}_p .

Lemma 7 For all matrices $A \in \mathcal{A}_p$, there holds:

- 1) Denote $B := A^{-1}$. Then, for all i, j $B_{ij} \geq 0$.

- 2) If $A_{ij} = 0$, then $\mathbf{s.e.i.}(A - h_i h_j^\top - h_j h_i^\top) \geq \mathbf{s.e.i.}(A)$. (Recall that h_i denotes i th canonical vector.)
 3) $\mathbf{s.e.i.}(A) \leq \frac{p}{2n-p+1}$, where the equality holds if and only if $A = A^* := 2nI - 11^\top + I$.

Proof: The claim 1) in Lemma 7 follows from the fact that the matrix A is an M matrix [30], and, consequently, the matrix $B = A^{-1}$ has all entries greater than or equal to zero (e.g., [30]). We now show the claim 2) in Lemma 7. Remark first that the function $\mathbf{s.e.i.}(\cdot)$ is convex and differentiable on the set of positive definite matrices. Applying the first order Taylor expansion lower bound at A , we get:

$$\mathbf{s.e.i.}(A - h_i h_j^\top - h_j h_i^\top) \geq \mathbf{s.e.i.}(A) + \text{tr} \left(\nabla \mathbf{s.e.i.}(A) \left(-h_i h_j^\top - h_j h_i^\top \right) \right), \quad (28)$$

where $\nabla \mathbf{s.e.i.}(A)$ stands for the (matrix form) gradient of $\mathbf{s.e.i.}(\cdot)$ at A and is equal to $-A^{-1} \mathbf{1}_p \mathbf{1}_p^\top A^{-1}$. Now, by claim 1), all entries of A^{-1} are nonnegative, and, thus, the second term on the right hand side of the inequality (28) is nonnegative as well. This completes the proof of claim 2.

We proceed and prove the claim 3) in Lemma 7. From claim 2) we know that the more -1 's a matrix $A \in \mathcal{A}_p$ has on its off-diagonal entries, the higher the value $\mathbf{s.e.i.}(A)$ can be. Therefore, A^* is a maximizer of $\mathbf{s.e.i.}$ over the set \mathcal{A}_p . Also, it is straightforward to check that $\mathbf{s.e.i.}(A^*) = \frac{p}{2n-p+1}$. We will show next that A^* is in fact the only maximizer of $\mathbf{s.e.i.}$ (over \mathcal{A}_p). To show this, it suffices to show that, for any choice of $1 \leq i, j \leq p, i \neq j$, the following strict inequality holds:

$$\mathbf{s.e.i.}(A^* + h_i h_j^\top + h_j h_i^\top) < \mathbf{s.e.i.}(A^*). \quad (29)$$

To this end represent the matrix $h_i h_j^\top + h_j h_i^\top$ as $h_i h_j^\top + h_j h_i^\top = HCH^\top$, where $H = [h_i \ h_j] \in \mathbb{R}^{n \times 2}$ and $C \in \mathbb{R}^{2 \times 2}$, $C_{12} = C_{21} = 1$, $C_{11} = C_{22} = 0$. Using the matrix inversion lemma, we get:

$$\mathbf{s.e.i.}(A^* + h_i h_j^\top + h_j h_i^\top) = \mathbf{1}_p^\top B^* \mathbf{1}_p - \mathbf{1}_p^\top B^* H \left(C^{-1} + H^\top B^* H \right)^{-1} H^\top B^* \mathbf{1}_p. \quad (30)$$

After straightforward algebra, we obtain:

$$\mathbf{1}_p^\top B^* H \left(C^{-1} + H^\top B^* H \right)^{-1} H^\top B^* \mathbf{1}_p = \frac{2n+1}{n+1 + \frac{1}{2n-p+1}}.$$

Since this term is greater than zero for all $1 \leq p \leq n$, the inequality (29) follows. This completes the proof of claim 3) and the proof of the Lemma. ■

We proceed with the proof of Theorem (1). The decision version of (4), for $k_1 = k_0 = +\infty$, is:

D-KL: Decision version of (4) “For given data: 1) vectors $m_0, m_1 \in \mathbb{R}^n$; 2) positive definite matrices S_0 and S_1 ; 3) positive integer $p, p \leq n$; and 4) a number f_{KL}^\bullet , determine whether there is a $n \times p$ sensor selection matrix E , such that $f_{\text{KL}}(E)$ defined in eqn. (6) is at least f_{KL}^\bullet .”

We now reduce the MQP to D-KL. Consider a simple, undirected graph \mathcal{G} and consider MQP of

determining whether there is a clique in \mathcal{G} of the size at least p . Define the matrix $S(\mathcal{G})$ as in eqn. (27). Consider an instance of D-KL, for some fixed p , with the following data: 1) $m_1 = \mathbf{1}_n$, $m_0 = \mathbf{0}_n$; 2) $S_1 = S_0 = S(\mathcal{G})$; 3) p ; and 4) $f_{\text{KL}}^\bullet = \frac{1}{2} \frac{p}{2n-p+1}$. Now, the answer to D-KL is YES (resp. NO) if and only if \mathcal{G} has (resp. does not have) a clique of size at least p . Thus, MQP problem is reduced to D-KL.

MQP problem reduces to the decision version of problem (5) (denoted by D-C), for $k_1 = k_0 = +\infty$, in a very similar way as it reduces to D-KL, by considering the instance of D-C with the data m_1, m_0, S_1, S_0, p , same as we considered for D-KL, but $f_{\text{KL}}^\bullet = \frac{1}{2} \frac{p}{2n-p+1}$ is replaced by $f_C^\bullet = \frac{1}{8} \frac{p}{2n-p+1}$.

B. Proof that the Kullback-Leibler and Chernoff distances are not submodular functions

See [23] for the definition of a submodular function. Consider two Gaussian distributions $\mathcal{N}(m_i, S_i)$, $i = 0, 1$, with parameters $m_0 = m_1$, $S_0 = I_3$ and $S_1 = I_3 + \epsilon(h_2 h_3^\top + h_3 h_2^\top)$ ($h_2, h_3 \in \mathbb{R}^3$), where $0 < \epsilon < 1$. Let $E_1 = h_1 \in \mathbb{R}^3$, $E_{13} = [h_1 \ h_3]$, $E_{12} = [h_1 \ h_2]$ and $E_{123} = I_3$.

Computing the Kullback Leibler distance for selections E_1, E_{12} and E_{13} we get $f_{\text{KL}}(E_1) = f_{\text{KL}}(E_{12}) = f_{\text{KL}}(E_{13}) = 0$, whereas $f_{\text{KL}}(E_{123}) = -\frac{1}{2} \log(1 - \epsilon^2) > 0$. Thus, $f_{\text{KL}}(E_{13}) - f_{\text{KL}}(E_1) < f_{\text{KL}}(E_{123}) - f_{\text{KL}}(E_{12})$ which proves that function f_{KL} is not submodular. The proof for the case of Chernoff distance can be done in similar way.

C. Solution to problem (25)

As explained in subsection IV-A1, the cost function can be written as $\frac{1}{2} \sum_{i=1}^p \phi_{\text{KL}}(\lambda_i(P^\top S P))$, where $\phi_{\text{KL}}(x) = x - \log x - 1$. Now, invoking Poincaré separation theorem, we can write (25) as:

$$\begin{aligned} & \text{maximize} \quad \frac{1}{2} \sum_{i=1}^p \phi_{\text{KL}}(x_i) \\ & \text{subject to} \quad \lambda_i(S) \leq x_i \leq \lambda_{n-p+i}(S), \quad i = 1, \dots, p \\ & \quad \quad \quad x_1 \leq x_2 \leq \dots \leq x_p \\ & \quad \quad \quad \exists P \in \mathbb{R}^{n \times p} \text{ s.t. } P^\top P = I_p \text{ and } x_i = \lambda_i(P^\top S P) \text{ for } i = 1, \dots, p. \end{aligned} \tag{31}$$

Next, we relax the problem (31) in the sense that we do not require that variables x_i are generated via the Stiefel matrix P :

$$\begin{aligned} & \text{maximize} \quad \frac{1}{2} \sum_{i=1}^p \phi_{\text{KL}}(x_i) \\ & \text{subject to} \quad \lambda_i(S) \leq x_i \leq \lambda_{n-p+i}(S), \quad i = 1, \dots, p \\ & \quad \quad \quad x_1 \leq x_2 \leq \dots \leq x_p. \end{aligned} \tag{32}$$

We now focus on problem (32). Function $\phi_{\text{KL}}(\cdot)$ is convex, with minimum equal to zero attained at point 1. Thus, its maximum over a compact interval is always at a boundary point. Also, if both boundary points are greater (resp. smaller) than 1, then the maximizer is the right (resp. left) boundary point. Therefore, if $\lambda_i(S) \geq 1$, then optimal x_i equals $\lambda_{n-p+i}(S) \geq 1$. Similarly, if $\lambda_{n-p+i}(S) \leq 1$,

optimal x_i equals $\lambda_i(S)$. If we rule out all eigenvalue intervals that fall in either of the two previous categories, we are left with intervals where 1 is in their interiors. We now focus on those eigenvalue intervals, i.e. on those $i = 1, \dots, p$ s.t. $\lambda_i(S) < 1$ and $\lambda_{n-p+i}(S) > 1$. It can be shown that there will exist a switching index $i_{\text{switch}} \in \{0, 1, \dots, p\}$ such that for all indices on the left of (or equal to) i_{switch} , optimal x_i is on the left boundary point, while for all indices on the right of i_{switch} optimal x_i is on the right boundary point. Summing up, a solution x^* of (32) will be always of the form $x^* = (\lambda_1(S), \dots, \lambda_{i_{\text{switch}}}(S), \lambda_{n-p+i_{\text{switch}}+1}(S), \dots, \lambda_n(S))^T$. (Case $i_{\text{switch}} = 0$ corresponds to $x^* = (\lambda_{n-p+1}(S), \dots, \lambda_n(S))^T$.)

We now construct solution of (31) from x^* -solution of (32); x^* consists of a subset of eigenvalues of the matrix S ; moreover, it is given by $x^* = (\lambda_1(S), \dots, \lambda_{i_{\text{switch}}}(S), \lambda_{n-p+i_{\text{switch}}+1}(S), \dots, \lambda_n(S))^T$, for some $i_{\text{switch}} \in \{0, 1, \dots, p\}$ (that must be determined). Therefore, there exists a matrix P^* that generates x^* , and hence solves (31). The columns of the matrix P^* are simply orthonormal eigenvectors of S corresponding to the eigenvalues $\lambda_1(S), \dots, \lambda_{i_{\text{switch}}}(S), \lambda_{n-p+i_{\text{switch}}+1}(S), \dots, \lambda_n(S)$.

We remark that the solution P^* of (25) is obtained in at most $p + 1$ steps. The number of steps is smaller when some of the interlacing eigenvalue intervals do not contain 1. If the number of such intervals is k , then the number of steps that our algorithm requires is exactly $p - k + 1$.

REFERENCES

- [1] L. Rueda and M. Herrera, "Linear dimensionality reduction by maximizing the Chernoff distance in the transformed space," *Pattern Recognition*, vol. 41, no. 10, pp. 3138–3152, October 2008.
- [2] L. L. Scharf, *Statistical signal processing: detection, estimation and time series analysis*. New York: Addison-Wesley, 1991.
- [3] D. Bajović, B. Sinopoli, and J. Xavier, "Sensor selection for hypothesis testing in wireless sensor networks: a Kullback-Leibler based approach," in *48th IEEE Conference on Decision and Control*, Shanghai, China, December 2009, pp. 1659 – 1664.
- [4] —, "Robust linear dimensionality reduction for hypothesis testing with application to sensor selection," in *47th Annual Allerton Conference on Communication, Control, and Computing*, Monticello, IL, September 2009, pp. 363 – 370.
- [5] C.-W. Ko, J. Lee, and M. Queyranne, "An exact algorithm for maximum entropy sampling," *Operations Research*, vol. 43, no. 4, pp. 684–691, July–August 1995.
- [6] E. L. Lawler and N. Draper, "Branch-and-bound methods: a survey," *Operations Research*, vol. 14, no. 4, pp. 699–719, July–August 1966.
- [7] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration for tracking applications," *IEEE Signal Processing Magazine*, vol. 19, no. 2, pp. 61 – 72, March 2002.
- [8] L. Kaplan, "Local node selection for localization in a distributed sensor network," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 42, no. 1, pp. 136–146, January 2006.
- [9] G. Hovland and B. McCarragher, "Dynamic sensor selection for robotic systems," in *IEEE International Conference on Robotics and Automation*, vol. 1, Albuquerque, NM, April 1997, pp. 272–277.

- [10] A. L. Buczak, H. H. Wang, H. Darabi, and M. A. Jafari, "Genetic algorithm convergence study for sensor network optimization," *Information Sciences*, vol. 133, pp. 267–282, 2001.
- [11] M. Chu, H. Haussecker, and F. Zhao, "Scalable information-driven sensor querying and routing for ad hoc heterogenous sensor networks," *International Journal of High Performance Computing Applications*, vol. 16, no. 3, pp. 293–313, August 2002.
- [12] F. Bian, D. Kempe, and R. Govindan, "Utility based sensor selection," in *IPSN '06, 5th International Conference on Information Processing in Sensor Networks*, New York, NY, USA, 2006, pp. 11–18.
- [13] E. Erten, J. Fisher, and L. Rotter, "Maximum mutual information principle for dynamic sensor query problems," in *IPSN '03, 2nd International Workshop on Information Processing in Sensor Networks*, Palo Alto, CA, April 2003, pp. 405–416.
- [14] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 451–462, February 2009.
- [15] V. Isler and R. Bajcsy, "The sensor selection problem for bounded uncertainty sensing models," *IEEE Transactions on Automation Science and Engineering*, vol. 3, no. 4, pp. 372–381, October 2006.
- [16] D. Bickson and D. Dolev, "Distributed sensor selection using a truncated Newton method," submitted for publication, available on arXiv:0907.0931v2 [cs.IT].
- [17] V. Kekatos and G. B. Giannakis, "Selecting reliable sensors via convex optimization," in *SPAWC 2010, IEEE International Workshop on Signal Processing Advances in Wireless Communications*, Marrakech, Morocco, June 2010.
- [18] Z. Song, C. R. Sastry, N. C. Tas, and Y. Q. Chen, "Feasibility analysis on optimal sensor selection in cyber-physical systems," in *ACC '09, American Control Conference*, St. Louis, MO, June 2009, pp. 5368–5373.
- [19] V. Fedorov, *Theory of optimal experiments*. New York: Academic, 1972.
- [20] Z. Quan, S. Cui, K. Ma, and A. H. Sayed, "Optimal linear fusion for distributed detection via semidefinite programming," *IEEE Transactions on Signal Processing*, vol. 58, no. 4, pp. 2431–2436, April 2010.
- [21] Z. Quan, S. Cui, and A. H. Sayed, "Optimal linear cooperation for spectrum sensing in cognitive radio networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 1, pp. 28–40, Feb. 2008.
- [22] M. K. Ali, L. Rueda, and M. Herrera, "On the performance of Chernoff-distance-based linear dimensionality reduction techniques," in *Canadian AI 2006, 19th Conference of the Canadian Society for Computational Studies of Intelligence*, vol. 4013, Quebec City, Canada, June 2006, pp. 3467–3478.
- [23] C. Guestrin, A. Krause, and A. P. Singh, "Near-optimal sensor placements in Gaussian processes," in *22nd International Conference on Machine learning*, vol. 119, Bonn, Germany, August 2005, pp. 265 – 272.
- [24] L. Brickman, "On the field of values of a matrix," *Proceedings of the American Mathematical Society*, vol. 12, no. 1, pp. 61–66, February 1961.
- [25] B. T. Polyak, "Convexity of quadratic transformations and its use in control and optimization," *Journal of Optimization Theory and Applications*, vol. 99, no. 3, pp. 553 – 583, December 1998.
- [26] Y. Nesterov and A. Nemirovskii, *Interior-point polynomial algorithms in convex programming*. Society for Industrial Mathematics, 1995.
- [27] B. Ricceri and S. Simons, *Minimax theory and applications*. Kluwer Academic Publishers, 1998.
- [28] R. A. Horn and C. R. Johnson, *Matrix Analysis*. New York: Cambridge University Press, 1985.
- [29] M. R. Garey and D. S. Johnson, *Computers and intractability: A guide to the theory of NP-completeness*. San Francisco: W. H. Freeman, 1979.
- [30] C.-Y. Zhang and Y.-T. Li, "On schur complement of block diagonally dominant matrices," *Linear Algebra and its Applications*, vol. 414, pp. 533–546, April 2006.